SJ Research
Maintenance Contracts

Version 1.01
September 1992

RESEARCH

SJ RESEARCH LIMITED
J1 THE PADDOCKS
347 CHERRY HINTON ROAD
CAMBRIDGE CB1 4DH

TELEPHONE 0223 416715
TECHNICAL SUPPORT 0223 414532
FAX 0223 416440

Disclaimer

While every care has been taken in the preparation of this document, we assume no responsibility for errors or omissions.
All goods and services provided by SJ Research Ltd. are subject to the company's Maintenance Agreement Terms of Contract. A copy of such terms is available on request and such standard terms shall apply on and from the date of actual commencement of such service or the provision of goods.

This edition (version 1.01) first published September 1992
Published by SJ Research Limited

Contents

Maintenance Contracts

SJ Research offers maintenance contracts on all its major items of equipment.

Why have a maintenance contract?

All computer equipment will eventually wear out and statistically a small percentage of units will fail prematurely. In the case of network servers a failure can affect a large number of users.
An empirically observed rule of computer equipment failure is that it always happens at the most inopportune moment and when you least expect it. A maintenance contact means that repairs and/or replacements are budgeted for and you should never have to find a large sum of money unexpectedly.

What repairs are covered?

A maintenance contract covers your equipment against any failure resulting from normal use of the equipment. It does not cover damage caused by accident, misuse or act of God. Damage by fire, flood or loss by theft should be covered by a general site insurance policy.

What can you cover?

Any product listed in the "Maintenance Contract Quotation Advice Note".

What must you cover?

A complete system. SJ Research systems are all designed in a modular fashion so that it is easy for you to upgrade your system. It is not possible to have a maintenance contract which covers only some of the items which make up a system. If the system fails then it may need to be examined as a whole in order to identify faults.

How much does it cost

The price of a maintenance contract clearly depends on exactly what equipment you have and how old it is. For this reason exact prices can only be given by quotation.

When can you take out a maintenance contract?

Obviously to ensure that you are covered at all times you should make sure that you take out a maintenance contract to run from the expiry of your warranty. However, it is possible to take out a maintenance contract on equipment at any time if you so wish.

Discounts for continuous cover

SJ Research operate a discount scheme for continuous cover of your equipment by one of our maintenance contracts.

Levels of cover

Class C - Extended warranty, preferential rates and support

- Over time it is likely that you will buy more equipment from SJ. We feel that it is important both for you and us to keep administration of maintenance contracts to a minimum. To achieve this we arrange for your maintenance contract always to fall due for renewal on the same date. You can add new pieces of equipment to it at any time and you will be charged on a pro-rata basis to bring it in line with your existing contract.

- SJ Research will organise and fund the collection and delivery by overnight carrier of units returned to SJ for repair.

- SJ Research will cover parts and labour charges for repair or total replacement with new equipment of equivalent or superior specification at SJ Research's sole discretion. If required we will attempt to recover data from corrupt discs. No guarantees given!

- SJ Research will provide automatic upgrades to documentation and software revisions*.

- SJ Research will provide a free first year subscription to the SJ User Group Newsletter.

- SJ Research will offer preferential rates for some of its other services.

**Software and documentation* revisions *are minor changes and bug fixes to existing release material. A software or documentation* release *may be a new chargeable product.*

Class A - Fast response extension

We will ensure that your system is on-line again in the fastest time possible. This will either be by sending you free of charge loan equipment via overnight carrier or by repairing your equipment on site, as deemed appropriate by SJ Research. If a fault is reported before 2.00pm then a replacement unit can be dispatched to arrive the next working day.

Class L - Labour charge only adjustment

For equipment over five years old the price of a contract is considerably reduced. This is because the cost of any parts required for repair is no longer covered. Labour, admin and carriage charges are still covered as normal.

Other SJ Research Services

SJ Research provides a comprehensive range of support services, and users under maintenance contracts are entitled to various discounts on these services.

System Support Contract

This is designed to provide you with general networking support for using Archimedes computers in a network environment. You are provided with access to the telephone hotline, application notes and software.

Training Courses

These are designed to tackle networking issues at all levels. They are held regularly in Cambridge.

Hire of equipment

It is normally possible to hire equipment from SJ Research. The minimum hire period is one week.

Technical Report

Line Drivers. Earth Potentials and Transient Suppressors

8th February 1990
A.R. Gordon, SJ Research

Line Drivers, Earth Potentials and Transient Suppressors

Many Econet users suffer repeated failures of line driver chips (75159 in BBCs, 26LS30/3691 in Masters) in their stations and servers. These failures are time-consuming and expensive to rectify, and in most cases are preventable by suitable precautions.

One of the weapons in the fight against line driver damage is the SJ Transient Suppressor Box (colloquially, "Zap box"). However, some thought is needed as to the best way to apply these devices.

This report describes in some detail the mechanisms involved in line driver failure, and the best ways to avoid them. Readers who are not interested in the technicalities should advance to the section headed "problem solving" where the conclusions are presented.

Why do line drivers fail?

In severe cases of electrical abuse, it is possible to damage any of the chips in the Econet interface, or indeed the whole computer. However, the line driver is the most vulnerable chip from the point of view of network-borne interference.

The reason for this is that the line driver is directly connected to the outside world. The other chips in the Econet interface are connected through large value resistores (or other chips), and these can absorb much of the energy of any interference entering through the Econet socket. However, the line driver has to supply significant power when it is transmitting across the network, and so has to have a direct connection to the network wires. Furthermore, the internal design of the chip is striving for the maximum possible current (in order to achieve high power output) rather than trying to minimise it for protection. The chip can, in fact, tolerate a short-circuit of its output, either to earth, or to the output of another chip, but an applied voltage above +12 volts, or below earth potential is just too much.

The manufacturers of the chip have provided some protection, in the form of diodes on the chip which prevent the voltage on the pins falling below earth potential, but all the energy entering the machine is dissipated as heat in a small area of the chip. If the overload is severe, or a small overload persists for any length of time, that part of the chip will melt and the chip will subsequently work poorly, if at all.

Given that each driver has two independant halves (for the two data lines in the cable), and overloads will tend to damage individual transistors, rather than the whole chip, it is possible to have "half blown" chips which are still working after a fashion, but with reduced performance. Damaged chips will either be "open circuit" so that no connection is made to the network - in this case, only that station is affected - or short circuit, permanently transmitting onto the network and disrupting all users. Since a different part of the chip comes into play when the station is powered off, it is possible to have faulty chips which work perfectly while the machine is switched on, but jam the network when it is switched off, or vice-versa.

In summary, the Econet interface does not permit applied voltages which are significantly above or below earth voltage. It will withstand overloads, but the amount of energy which can be absorbed is small.

Sources of interference

There are two major sources of damaging interference - externally induced pulses (thunderstorms etc.), and the electrical supply system. We do not expect the network to survive a direct lightning strike (which will probably set the building on fire anyhow), but we would like to tolerate the case of a bolt of lightning landing nearby.

When lightning strikes, there is a huge electrical current (many thousands of amps) entering the ground in a small area. This has two effects: there is a strong magnetic field, which will induce currents in any nearby conductors (such as the Econet cable), and there will also be potential differences in the soil, as the current spreads out into the general mass of the earth, such that buildings near the point of the lightning strike will measure a different earth potential to those further away. In either case, significant voltages will appear at the end of the cable, measured with respect to the local earth: either the voltage is induced down the length of the cable, or the cable is simply connecting two buildings that are at different potentials. A current flows from the earth in one building, through the line drivers in local stations, down the cable, in through the line drivers of machines in the second building and away to that earth.

Lightning problems generally only affect sites with more than one building - the wiring in a single building is not usually laid out in such a way as to couple the magnetic flux, and a single building will generally be at a fairly consistent earth potential. Electrical supply problems, however, can affect any size of network, although small networks will only be affected if there is a blatent defect.

In considering electrical supply problems, we are mainly concerned with the earthing arrangements. There will be high-voltage transients in the supply itself, but these will usually be adequately damped by the computer's power supply tranformer etc. and are unlikely to be a cause of problems; fitting mains transient suppressors to your 13A sockets gives little benefit.

As we have said, the voltages applied to a station should not be significantly above or below earth potential. The network cable connects the earths of all the computers together, so under normal conditions all the stations see the same earth potential and all is well. However, the network earth connection is reasonably light-weight, and cannot hope to iron out major differences in earth potentials that already exist. If the earth potential at one station is significantly above or below that at other stations, current will again flow from one earth to the other, through the line-driver chips in the respective machines.

What can cause earth potential differences? The most common situation is that the earth pin on a mains socket has become disconnected from the supply earth. In this case, the earth potential at that machine is completely free to float to any value, until the machine is connected to the network. However, if the computer has a switched-mode power supply (most BBCs do), there will usually be suppression components linking the live, neutral and earth connections; with earth disconnected, about 60V AC is observed on the "earth" line. If such a machine has been floating and is then connected to the network, the network earth will gradually bring it back into line, but not before the 60V has had time to damage a few line drivers. If there are several such machines, the network earth will not be able to cope with the leakage currents and there will be a permanent potential difference, probably sufficient to make the network work unreliably and to blow line drivers on a regular basis.

If the electrical supply wiring is in good condition and all machines are firmly connected to the system earth, there should be no significant potential difference between the machines in one building. However, separate buildings (or, exceptionally, buildings with more than one independant electrical supply system) do not have a direct connection between their earths, other than the network cable. The electrical supply systems will have earth electrodes driven into the ground, and providing

these are in good condition and the soil has not dried out, the earth potentials will be close enough for the network to operate without trouble. Occasionally, however, due to a faulty appliance being connected or simply switch-on surges in large plant, the earth potential will take a temporary jump, causing damage to line drivers unless precautions have been taken.

An additional cause of occasional problems is faulty appliances being plugged in near to the computers. A fault shorting live to earth will cause large currents to flow in the earth wiring until such time as the fuse blows. The wiring will have been designed to handle these loads and keep the potentials below the level that could cause fatal electric shocks, but line drivers are more sensitive than humans and may well fail.

What can transient suppressors do?

Inside a transient suppressor box, there are four suppressor diodes connecting between each of the network clock and data lines and ground. Under normal conditions, very little current flows through the diodes, but as soon as the voltage across the diodes exceeds specification, the diodes begin to conduct and "clamp" the voltage to a safe level.

The suppressor diodes are similar to conventional zener diodes, but have a much greater capacity to absorb energy. For example, the diodes used in SJ Research transient suppressors will keep the voltage below 13V even if a current of up to 130 Amps is flowing. They can absorb a 1kW of power for up to 5ms, and even higher power levels for shorter periods. Even if the energy in the surge exceeds the capacity of the diodes, they are designed to fail to a short circuit - damaging themselves, but continuing to protect the computers. By contrast, a normal zener diode will fail at very much lower energy levels, and will probably fail open-circuit, leaving the remainder of the surge energy to dissipate itself in the computers.

The purpose of a transient suppressor is to clamp each of the signal lines within a permitted range from its earth connection. If the earth potential at a transient suppressor is the same as that at a nearby station, this means that the voltages applied to that station will never be high enough to damage it.

Note that surge voltages do not appear instantaneously on the cable; they travel as pulses along the cable at the same speed as other signals. This means that if the pulse reaches a station before it reaches the transient suppressor, it will have damaged that machine before the transient suppressor has had time to act. Transient suppressors should therefore be installed between likely sources of interference an the machines that they are intended to protect; this usually means the point at which the network cable enters the building.

Problem Solving

To avoid basic problems with electrical supply:

Test the earth on all mains sockets before use with computers. Do this with all computers disconnected from the network, as the network earth can give the false impression that the mains earth is OK.

Avoid the use of 4-way mains leads. The flexible nature of the rubber housing means that poor contacts develop easily. Ideally, the building's fixed wiring should provide enough sockets, but if not, multi-way distribution boards can be made from standard mains sockets screwed to a block of

wood (or to table top etc), which will be much more reliable.

Beware of elderly electrical installations where there is no wire forming the earth connection, which simply relies on a metal conduit to provide earth. These often develop faults.

To counter problems due to faulty appliances being plugged in nearby, it is preferable to have separate mains circuits supplying the computers, which are not used for other appliances. This is probably too expensive to arrange in all circumstances, but in a computer room it may be appropriate to have one circuit for computers and a separate supply for vacuum cleaners etc.

Where there is more than one building (or electrical supply):

Each group of computers (eg. those in one building) should be protected by transient suppressors. The suppressors should be fitted in line with the cable where it enters the area containing the computers; possibly where the cable enters the building, but if all the computers are located close together it may be better to fit the suppressor where the cable enters the first room containing computers. The earth terminal on the suppressor box should be connected by a heavy-duty cable to the earth of the electrical supply serving the computers. Ideally, this should be the main earthing terminal on the electrical supply distribution board serving all the computers. If this is inaccessible, the earth of a socket where computers are to be used may serve as (an inferior) substitute. Taking the earth from sockets not associated with the computers is undesirable, as this will feed interference into the network from faulty appliances which may be plugged into that socket.

Where there are independant electrical supplies in different parts of a building, or no obvious connection can be found between supplies, the computers fed from each supply should be protected separately. Different phases from a three-phase supply do not matter, however, so long as all the circuits share the same earth.

A case was discovered recently where the earth for a transient suppressor had been taken from the building's lightning conductor. This is NOT advised!

These steps should protect computers from most transient variations in earth potentials. However, in some cases there may be a permanent difference in earth potentials, sufficient to permanently energise the suppressors. In this case, damage is unlikely to occur but the network will not work; in extreme cases, sparks may fly as the network earth is connected between the buildings!

Most permanent potential differences are due to poor earthing of the mains electrical supply. It may be possible to alleviate this by connecting the network earth to stakes driven into the ground, but this is not in general a good idea. If the original problem was caused by a corroded or dried-out earth stake on the electrical supply system, connecting a good stake to the network earth means that any earth fault current from the building electrical supply wil now travel down the network earth and through the new stake. While this may resolve day-to-day problems, in the event of a serious electrical fault in that building, far more damage will probably occur; in the worst case, it could be a fire hazard. If you believe your electrical supply earth is faulty, fix it in the electrical system, not in the network.

There may be rare cases where a permanent potential differnce exists legitimately. This may particularly be the case where the electrical supply is earthed by PME (where earth and neutral are connected together as the supply enters the building and the electricity board undertake to earth the neutral line), although it is very rare for such potentials to be large enough to cause problems. A particular case arises on very large sites where buildings are fed from different electrical sub-stations. In this case, fuses should be inserted in all network cables (including the earth) for electrical safety.

Locating faulty line driver

It is not usually practical to take precautions against absolutely all possible sources of interference, and so line drivers will very occasionally fail. Given that there can be faulty line driver chips affecting the whole network, how do we go about locating the faulty machines? Indeed, if the network is working moderately well, how do you determine if there are any faulty drivers at all?

The most useful check is to examine the voltages on the data lines when the network is running, but no users are active. Measuring the voltage between each data line and earth (with a simple voltmeter or oscilloscope), there should be about 2.2V on data+ and 1.8V on data- (these figures are approximate; D+ may be anything from 1.8 to 2.3 volts, and D- should be 0.3 to 0.5 volt less than D+).

If you believe that there is a problem, the first step is to disconnect all stations from the network, and ensure that all is well with just clock, network cable, suppressors and terminators. If even this does not work, remember that the clock box also contains a line driver, which is equally vulnerable. If the voltages are too high with no stations connected, suspect the network earth connection between the clock box and terminators.

When the empty network is working, re-connect the stations one by one, while continuing to monitor the data line voltages. Try each machine both switched on and switched off, to check for part-damaged chips. If the data line voltages change significantly when a machine is added, that machine almost certainly has a faulty line driver (or more serious fault).

Faulty line drivers in early BBCs (issue 3 or earlier circuit board) can also cause "no clock" on other stations, as they have a (totally spurious) connection from the line driver to the clock lines. This is removed in later BBCs.

When all machines are connected, check that each can communicate with the fileserver. Machines reporting Net Error or Station Not Present probably have totally blown line driver chips.

It is also possible to damage chips in such a way that they continue to work perfectly, but their internal connections have been partly destroyed, making them very sensitive to further mild overloads. It is therefore not unusual after a serious problem has been corrected for other line drivers to fail subsequently, for no apparent reason - most of the damage having occurred in the original event. Unfortunately, there is no practical way to identify such chips.